PHOTOGRAPH THIS SHEET

AFWAL-TR-87-1162

DOCUMENT IDENTIFICATION

Dec 1987

DISTRIBUTION STATEMENT

QUALITY
INSPECTED
2

DATE ACCESSIONED

DATE RETURNED

05 106

DATE RECEIVED IN DTIC

REGISTERED OR CERTIFIED NO.

PHOTOGRAPH THIS SHEET AND RETURN TO DTIC-DDAC

AFWAL-TR-87-1162

A METHODOLOGY FOR HARDWARE VERIFICATION BASED ON
LOGIC SIMULATION

Randal E. Bryant

Carnegie-Mellon University
Computer Science Department
Pittsburgh, PA   15213-3890

December 1987

Interim

Approved for Public Release; Distribution is Unlimited

AVIONICS LABORATORY
AIR FORCE WRIGHT AERONAUTICAL LABORATORIES
AIR FORCE SYSTEMS COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OHIO   45433-6543

NOTICE

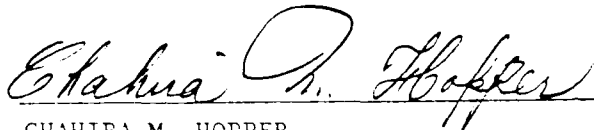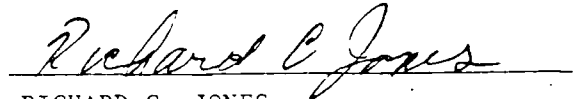When Government drawings, specifications, or other data are used for any purpose other than in connection with a definitely Government-related procurement, the United States Government incurs no responsibility or any obligation whatsoever. The fact that the Government may have formulated or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication, or otherwise in any manner construed, as licensing the holder, or any other person or corporation; or as conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

This report has been reviewed by the Office of Public Affairs (ASD/PA) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication.

CHAHIRA M. HOPPER
Project Engineer

RICHARD C. JONES.
Ch, Advanced Systems Research Gp
Information Processing Technology Br

FOR THE COMMANDER

EDWARD L. GLIATTI
Ch, Information Processing Technology Br
Systems Avionics Div

If your address has changed, if you wish to be removed from our mailing list, or if the addressee is no longer employed by your organization please notify AFWAL/AAAT , Wright-Patterson AFB, OH 45433-6543 to help us maintain a current mailing list.

Copies of this report should not be returned unless return is required by security considerations, contractual obligations, or notice on a specific document.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No 0704-0188

1a REPORT SECURITY CLASSIFICATION Unclassified	1b RESTRICTIVE MARKINGS
2a SECURITY CLASSIFICATION AUTHORITY	3 DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.
2b DECLASSIFICATION DOWNGRADING SCHEDULE	

4 PERFORMING ORGANIZATION REPORT NUMBER(S) CMU-CS-87-128	5 MONITORING ORGANIZATION REPORT NUMBER(S) AFWAL-TR-87-1162

6a NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University	6b OFFICE SYMBOL (If applicable)	7a NAME OF MONITORING ORGANIZATION Air Force Wright Aeronautical Laboratories AFWAL/AAAT-3
6c ADDRESS (City, State, and ZIP Code) Computer Science Dept Pittsburgh PA 15213-3890		7b ADDRESS (City, State, and ZIP Code) Wright-Patterson AFB OH 45433-6543

8a NAME OF FUNDING SPONSORING ORGANIZATION	8b OFFICE SYMBOL (If applicable)	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F33615-84-K-1520
8c ADDRESS (City, State and ZIP Code)		10 SOURCE OF FUNDING NUMBERS

PROGRAM ELEMENT NO	PROJECT NO	TASK NO	WORK UNIT ACCESSION NO
61101E	4976	00	01

11 TITLE (Include Security Classification)

A Methodology for Hardware Verification Based on Logic Simulation

12 PERSONAL AUTHOR(S)
Randal E. Bryant

13a TYPE OF REPORT Interim	13b TIME COVERED FROM _____ TO _____	14 DATE OF REPORT (Year, Month Day) 1987 December	15 PAGE COUNT 38

16 SUPPLEMENTARY NOTATION

17	COSATI CODES		18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	

19 ABSTRACT (Continue on reverse if necessary and identify by block number)

A logic simulator can prove the correctness of a digital circuit if it can be shown that only circuits implementing the system specification will produce a particular response to a sequence of simulation commands. This style of verification has advantages over other proof methods in being readily automated and requiring less attention to the low-level details of the design. It has advantages over other approaches to simulation in providing more reliable results, often at a comparable cost.

This paper presents the theoretical foundations of several related approaches to circuit verification based on logic simulation. These approaches exploit the three-valued modeling capability found in most logic simulators, where the third value X indicates a signal with unknown digital value. Although the circuit verification problem is NP-hard as measured in the size of the circuit description, several techniques can reduce the simulation complexity to a manageable level for many practical circuits.

20 DISTRIBUTION/AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ DTIC USERS	21 ABSTRACT SECURITY CLASSIFICATION Unclassified
22a NAME OF RESPONSIBLE INDIVIDUAL Chahira M. Hopper	22b TELEPHONE (Include Area Code) (513) 255-7865	22c OFFICE SYMBOL AFWAL/AAAT-3

DD Form 1473, JUN 86 Previous editions are obsolete.

A Methodology for Hardware Verification Based on Logic Simulation*

Randal E. Bryant
Computer Science Department
Carnegie Mellon University
Pittsburgh, Pennsylvania 15217

June 8, 1987

Abstract

A logic simulator can prove the correctness of a digital circuit if it can be shown that only circuits implementing the system specification will produce a particular response to a sequence of simulation commands. This style of verification has advantages over other proof methods in being readily automated and requiring less attention to the low-level details of the design. It has advantages over other approaches to simulation in providing more reliable results, often at a comparable cost.

This paper presents the theoretical foundations of several related approaches to circuit verification based on logic simulation. These approaches exploit the three-valued modeling capability found in most logic simulators, where the third value X indicates a signal with unknown digital value. Although the circuit verification problem is NP-hard as measured in the size of the circuit description, several techniques can reduce the simulation complexity to a manageable level for many practical circuits.

1   Introduction

Logic simulators provide a valuable tool for testing the correctness of digital circuits. Typically, however, only a limited set of test cases is simulated, and the circuit is presumed correct if the simulator yields the expected results for all cases. Unfortunately, this form of simulation leaves the designer uncertain that all circuit design errors have been eliminated.

Stories abound of errors that remain undetected despite many hours of simulation and even actual circuit operation. Conventional wisdom holds that logic simulators are incapable of more rigorous verification. They are viewed in the same class as program debuggers— useful tools for informal testing, but nothing more.

Formal verification involves proving that, under some abstract model of system operation, the circuit will behave as specified for all possible input sequences. A formal proof gives strong confidence that the circuit will function correctly. In this paper, we will show that a logic simulator can form the basis of a formal verifier. At first, this claim might seem both obvious and of little practical value, since most systems are too complex to simulate exhaustively. We will argue to the contrary on both points. When the circuit has potential for sequential behavior, even simulating all possible input patterns may fail to detect an error. Furthermore, verification by simulation can be made practical for a significant class of circuits.

Formal verification does not guarantee that the actual circuit will operate properly. The assumptions made in the abstract model may not hold in the physical implementation. For example, most methods of verifying digital systems assume that the circuit adheres to a logic abstraction whereby all signals can be represented by discrete values. Without such an abstraction, verification would be tedious, if not impossible. Design errors that cause marginal, nondigital circuit behavior may not be detected by verification against such a model. Similar problems arise in program verification. For example, most proofs of program correctness abstract the finite arithmetic implemented by computers as operations over the integer or real domain. A verification against such a model cannot detect errors due to arithmetic overflow or underflow. In discussing formal verification, we must remember that the level of confidence it provides is only as strong as the degree to which the abstract model matches actual system operation.

1.1   Structural Approaches

Most hardware verification methodologies [1,2,16,17,21,22,23,24] utilize *structural* techniques. In such an approach, the circuit is described hierarchically, where a component is defined at one level in the hierarchy as an interconnection of components defined at lower levels. The system specification consists of a description of the behavior of every component at each level of the hierarchy. Verification then involves proving that each component implements its part of the specification, assuming that its constituent components implement theirs.

Structural verifiers have several noteworthy strengths. They can exploit the circuit hierarchy to reduce proof complexity, since a proof is required only for each unique circuit component. Many large, but highly structured circuits have been verified structurally. Second, they can naturally be extended to parameterized circuit descriptions, proving the

correctness of entire families of circuits [9]. Finally, structural verifiers can apply different modeling abstractions according to the level in the hierarchy, such as representing signals at lower levels as bits and at higher levels as integers [2].

On the other hand, these verifiers have several shortcomings. Even when automated, they require the user to specify the intended behavior of each component in the circuit hierarchy. The verifier serves largely as a "proof checker", making sure that each component fulfills its specification. Many circuits are not designed to facilitate component specifications, and hence verification requires much tedious effort on the part of the user. Consider, for example, an adder circuit that utilizes carry-lookahead. Although the addition function is straightforward to specify, the low level details of the implementation are complex. Furthermore, the circuit contains many different component types and hence requires a lengthy specification and verification. For such a circuit, a verification method that allows the user to deal with the overall input-output behavior would be far preferable.

As a second shortcoming, most structural verifiers use highly simplified models of electrical and timing behavior to make the proof and component specifications tractable. Most assume, for instance, that the circuit components operate as unidirectional logic elements computing outputs in response to their inputs. In actual circuits electrical behavior can be far more subtle, such that the behavior of a component depends on its operating environment. As an example, the direction of information flow through a CMOS transmission gate is determined solely by the driving capabilities of the circuitry at either end [10]. Clearly, any specification of such a gate must include restrictions on the environment in which it is placed. As a notable exception to these highly simplified models, Weise [23,24] has developed a verifier that proves the correctness of MOS circuits under a model that includes detailed electrical and timing information. His verifier automatically checks every environment in which components are placed for compliance with the preconditions for correct operation. In general, however, prospects do not look good for automating structural verifiers to the point where circuits can be verified with little manual effort and with realistic circuit models. Formulating the proper set of assertions about each component requires a more sophisticated reasoning capability than will be automated in the near future.

1.2   A Behavioral Approach

This paper proposes a *behavioral* approach to circuit verification. In this approach, the verifier applies logic simulation to compute the circuit response to a series of stimuli chosen to detect all possible design errors. The user is freed from the tedium of proving the correctness of every component. Instead, the circuit is viewed at a higher level in terms of its desired input-output or state transition behavior. More realistic circuit models can be used, because only the simulator need be concerned with the modeling details.

3

Although this approach to hardware verification overcomes several weaknesses of structural verifiers, it cannot match some of their strengths. Simulation cannot exploit hierarchy very effectively, because the different instances of a component can have different stimuli and hence must all be evaluated. There is also no known way to simulate an entire class of circuits in a single run. Perhaps the ideal verifier would combine both styles. A hybrid approach would use behavioral verification to prove the correctness of a set of components forming some intermediate level in the circuit hierarchy. This would avoid the need to specify the behavior of the low level components and could employ the detailed circuit models required at these levels. Structural verification would then be applied to the hierarchical composition of the intermediate components, exploiting the regularity of their interconnections and their higher abstraction levels. Thus the work presented here should be viewed as complementing structural verification, rather than seeking to replace it.

1.3 Overview of the Methodology

The task of evaluating a circuit by simulating its response to a set of stimuli relates closely to the "machine identification" problem first described by Moore [18]. He showed that, in general, no finite set of stimuli could fully characterize the behavior of a sequential system. He suggested overcoming this problem by fixing an upper bound on the total number of system states. Unfortunately, for circuits of significant size, this bound is too high for Moore's identification algorithm to be practical. Instead, our method overcomes the identification problem by simulating the behavior over a three-valued domain with conventional Boolean values 0 and 1, plus a value X representing an undefined or uninitialized signal. Such a capability is found in most logic simulators [12]. In addition to supplying input patterns during simulation, we assume that the user can issue ERASE commands, causing all state variables to be set to X. Although the power of three-valued simulation has been studied extensively in the context of hazard detection [6,15,25], its potential role in circuit verification has not been widely recognized.

In the interest of generality and simplicity, the paper views hardware specification, circuit behavior, and logic simulation in a rather abstract way. The desired behavior is specified by a (Moore model) finite state automaton. The circuit is also a finite state automaton, with a particular binary coding of the states. Although this is an unconventional view of a circuit, we will argue its appropriateness for behavioral verification, where the focus is on how the circuit operates rather than how it is constructed. Circuit verification involves proving that the specification and circuit automata have equivalent input-output behavior. The simulator models the behavior of the circuit automaton, computing new state and output values in response to inputs supplied by the user. A mild, monotonicity property is imposed on the simulation of three-valued behavior to capture the notion that X represents an unknown or ambiguous digital value.

The style of simulation required to prove correctness depends on the nature of the system

specification. A *definite* system [14,19], for which the behavior depends on only a bounded number of previous inputs, can be verified by straightforward "black-box" simulation. Black-box simulation involves simply observing the output produced by the simulated circuit in response to a sequence of input and ERASE commands with no consideration of the internal circuit structure. Verifying the implementation of an *indefinite* system, on the other hand, requires a more implementation-specific "state transition" simulation. With this method, key circuit state variables are identified, and the different possible state transitions simulated. In either case, the verification requires little, if any, understanding of the detailed circuit design.

Circuit verifiers can err in two different ways. One that rejects a correct circuit gives a *false negative* response, whereas one that accepts an incorrect circuit gives a *false positive* response. This paper is concerned mainly with avoiding false positive responses. Such a response has more potential for danger—it may cause a defective design to be implemented or put into service. Furthermore, deciding whether a simulator has produced a false negative response requires more detailed information about the circuit electronics and the simulation algorithm than can be presented in a general way. However, for the simulation sequences presented in this paper, a false negative response must have a particular form, namely the simulator will produce X on some output when 0 or 1 was expected.

As mentioned earlier, any approach to formal verification guarantees proper circuit operation only if the assumptions made in the abstract model hold in the circuit implementation. For the case of simulation-based verification, we must assume that the actual circuit behaves identically to its simulation. When a circuit has been "verified" by simulation, it simply means that it has no errors that could be detected by simulating additional patterns. It is important to maintain this perspective on the problem addressed by this paper. It reflects a weakness intrinsic to any approach to verification and not to simulation alone.

1.4   Contents of Paper

This paper presents both theoretical and practical aspects of a hardware verification methodology based on multi-valued logic simulation. Section 2 illustrates the key ideas by means of several circuit examples. Section 3 gives a notation and mathematical background for describing system specifications, digital circuits, logic simulators, and the verification problem. Section 4 gives a formal characterization of the capabilities and limitations of black-box simulation. Section 5 shows how these limitations are overcome by state transition simulation.

Section 6 discusses methods to improve the computational efficiency of the verifier. Two methods are proposed to reduce the computational effort in verifying large digital systems: input weakening and symbolic simulation. Input weakening involves using the value X on an input to represent "don't care", thereby reducing the total number of patterns that need

to be simulated. Symbolic simulation involves augmenting a simulator with a symbolic Boolean manipulator to compute the behavior of the circuit over input patterns containing Boolean variables. Section 7 demonstrates a practical application of the methodology to the verification of a random-access memory. By exploiting input weakening, an n-bit memory can be fully verified by simulating $O(n \log n)$ patterns, even though the circuit has 2^n possible states. Section 8 concludes the paper with a discussion of the method.

All circuit examples shown in this paper are designed in MOS technology. This choice reflects the historical background of the research as well as the belief that MOS circuits form a particularly difficult class for hardware verification. In particular, state can be stored as dynamic charge on capacitive nodes. Unlike circuits where all state is stored in feedback loops, there is no straightforward way to identify all state variables. In fact, design errors commonly introduce unintended state variables and sequential dependencies. The methodology presented here, however, applies to most digital technologies, with the caveat that circuits are assumed to operate on synchronized input data. Asynchronous systems seems to call for more powerful class of verification tools, such as the model checker of Clarke, *et al* [3,7], since they cannot be viewed simply as processing a single sequence of input data.

This research provides two major contributions to the state of the art in circuit validation. First, it presents a simple, yet powerful, method of proving the correctness of digital hardware. Second, it provides insights into ways to better utilize a simulator even when only informal validation is sought. It shows that by exploiting a latent capability found in most logic simulators, namely three-valued modeling, more rigorous validation can be obtained at comparable cost.

2 Verification Examples

Before proceeding with the mathematical formalism, we present the main concepts via several circuit examples. These examples illustrate several pitfalls of simulation and how three-valued modeling can be exploited to overcome them.

2.1 Definite Systems

Consider the seemingly simple task of proving that a circuit implements a NOR logic gate. Figure 1 shows two proposed implementations in CMOS technology [10]. If we were to simulate these circuits using a simulator that can model a MOS circuit at the transistor level[4], the following responses would be produced when the input patterns are applied in the sequence shown:
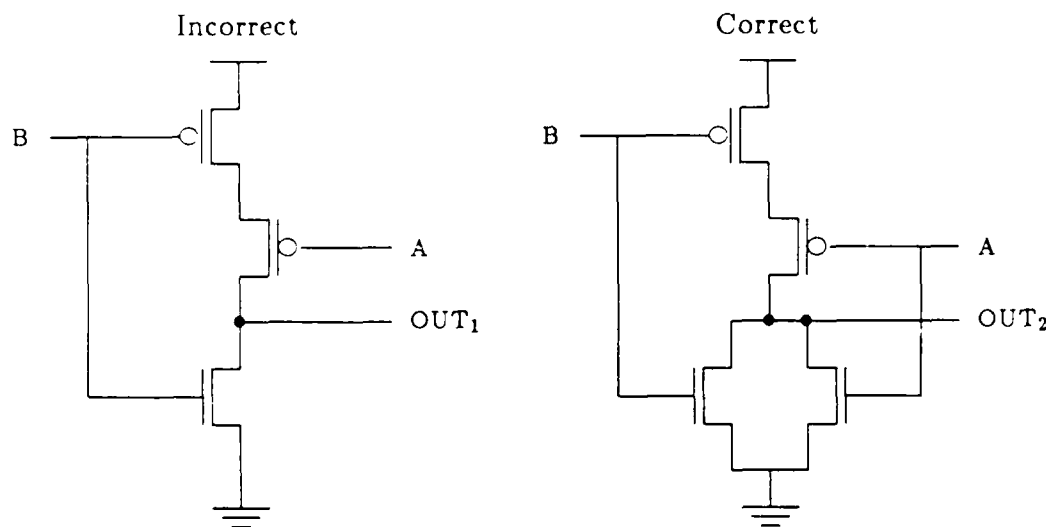
Figure 1: Implementations of a NOR Gate in CMOS

A	B	OUT$_1$	OUT$_2$
0	0	1	1
0	1	0	0
1	1	0	0
1	0	0	0

The two circuits appear identical for all possible input combinations, as they would be in actual implementations of the circuits. However, only the second circuit is a valid NOR gate. The first is a two-state sequential circuit, because when A = 1 and B = 0, the output node is electrically isolated from all others and remains charged at its previous value. Due to the order in which the input combinations were applied, it just happened that the previous value of OUT$_1$ equaled the value a NOR gate should produce for this input combination. On the other hand, had this input combination been simulated immediately after the combination A = B = 0, the output would have been 1.

This example illustrates a common problem in testing a circuit by simulation even when rigorous verification is not sought. A design error introduces an unintended sequential dependency in the circuit, but this error remains undetected because of the particular order in which the test sequences are simulated. Clearly, such a condition is not acceptable for formal verification.

Suppose, on the other hand, that an ERASE command is given before simulating each input combination, causing all state variables to be set to X. Such a simulation would produce the following results:

A	B	OUT$_1$	OUT$_2$
0	0	1	1
0	1	0	0
1	1	0	0
1	0	X	0

The presence of an X on OUT$_1$ indicates that the output of this circuit may not be uniquely defined when A = 1 and B = 0 (or it could be a false negative response by a valid circuit). On the other hand, each input combination produces a unique response for the second circuit, and since the responses match those of a NOR gate, one can conclude the circuit is correct. This conclusion can be drawn without any further information about the structure of the circuit or the number of state variables.

A combinational system such as a NOR gate is 1-definite; its output at any time depends only on the most recent input. The method shown above generalizes to any definite system specification, where the output depends only on the most recent k inputs for some constant k. That is, suppose for every possible input sequence of length k, setting all state variables to X and then simulating the sequence yields an output equal to the desired value. One can then safely conclude the circuit implements the specification.

Observe that this requirement for definiteness applies only to the system specification and not to the circuit being evaluated. For example, the simulator is able to detect the incorrect NOR gate even though the circuit itself is not definite. In particular, its output depends on inputs that occurred arbitrarily long ago as long as A is held at 1 and B at 0.

2.2   Indefinite Systems

Many sequential systems are not definite. For example, a simple 1-bit latch has an output dependent on an input that occurred arbitrarily long in the past as long as no new value is written into it (by setting its LOAD input to 1). For such a system, given any value k, there will always be an input sequence of length k that does not cause the system to produce a unique output. When simulating this sequence following an ERASE command, even a correctly designed circuit will give an X on the output.

In Section 4 we will show that for any indefinite system, there is no way to prove that a circuit implements its specification by simply observing the output values resulting from a sequence of input patterns and ERASE commands. This general limitation of black-box simulation can be illustrated using a 1-bit latch circuit. Consider a simulation sequence that is claimed to detect any defective latch design. Since the sequence is finite, there must be some value l such that the LOAD input is never held at 0 for l or more consecutive patterns.

8

Figure 2: Latch with Booby Trap

Consider the circuit of Figure 2 consisting of a correct latch with additional circuitry implementing a "booby trap". Whenever a value is written into the latch, all bits in the booby trap shift register are cleared to 0. The output of the shift register is EXCLUSIVE-OR'ed with the latch output to produce the circuit output OUT′. Thus, the circuit behaves as a proper latch as long as the shift register output equals 0. If no further data is written, a 1 will shift through the register until it ultimately causes the circuit output OUT′ to be complemented. Clearly, this circuit does not behave as a latch should. However, the proposed sequence does not cause enough consecutive shift operations for the defective circuit to behave differently from the correct one.

Less obviously, even attempts to expose the booby trap by giving ERASE commands or by giving input sequences containing X's will fail to distinguish the correct latch from the circuit of Figure 2. Unlike the NOR gate example, any action that would cause the simulator to produce output X for the incorrect circuit would also cause it to produce output X for the good circuit. Thus, the proposed simulation sequence cannot distinguish between a correct circuit and this incorrect one. This argument holds for any simulation sequence by making l sufficiently large.

To verify indefinite systems, more information is required about the circuit state variables and their relation to the states of the system specification. However, in the spirit of black-box simulation, we would like to minimize the amount of detail about the circuit structure that the user must provide. To achieve this goal, assertions about the state variables and how they are transformed by the input values are expressed in a notation similar to the Floyd-Hoare assertion method of program verification [8,11]. Each assertion is then verified by a short simulation sequence.

Figure 3: Latch Circuit

A *circuit assertion*, denoted by an equation of the form $P \{ I \} N \wedge O$, consists of 4 predicates over the circuit input, output, and state variables. Predicate P specifies a precondition on the state variables, I an action on the input variables, N a postcondition on the state variables, and O a postcondition on the output variables. Each predicate is the logical conjunction of terms of the form $v = 1$ or $v = 0$, where v is circuit variable. A circuit assertion can be interpreted as a statement that for any circuit state satisfying P and for any action satisfying I, the new circuit state will satisfy N, and the output will satisfy O.

As an example, consider an implementation of the 1-bit latch illustrated in Figure 3. No internal details of the circuit are shown except that the bit is stored in a feedback path containing two electrical nodes Q and $\overline{Q}$. The following assertions specify the state transition behavior of the circuit where a ranges over 0 and 1:

$$\textit{true} \left\{ \text{DATA} = a \wedge \text{LOAD} = 1 \right\} Q = a \wedge \overline{Q} = \neg a \wedge \text{OUT} = a$$

$$Q = a \wedge \overline{Q} = \neg a \left\{ \text{LOAD} = 0 \right\} Q = a \wedge \overline{Q} = \neg a \wedge \text{OUT} = a$$

The first equation asserts that a write operation sets the state of a latch, while the second asserts that the latch state does not change as long as no new data is written. In both cases the value of OUT should equal that of Q.

Given an assertion $P \{ I \} N \wedge O$, its verification by simulation involves the following steps. An ERASE command is given to set all state variables to X. For each term $v = a$ in P (respectively I), state variable v (resp. input variable v) is set to a. All input variables not occurring in I are set to X. The simulator then computes the resulting output and new state. For each term $v = a$ in N (respectively O), state variable v (resp. output variable v) is tested for equality with a. If all of these tests hold, then the assertion is proved.

For example, consider the latch circuit of Figure 3, and the faulty circuit formed by adding the booby trap of Figure 2 to it. Simulating the four sequences specified by the assertions

would yield the following results:

Initial Values				Results			
Q	$\overline{Q}$	DATA	LOAD	Q	$\overline{Q}$	OUT	OUT$'$
X	X	0	1	0	1	0	0
X	X	1	1	1	0	1	1
0	1	X	0	0	1	0	X
1	0	X	0	1	0	1	X

The defective circuit passes the first two tests when new data is written, because this causes the shift register to be cleared. For the final two tests, however, all state variables with the exception of Q and $\overline{Q}$ are initialized to X. The shift register output will remain at X, causing an X to appear on OUT$'$ and the tests to fail.

3 Mathematical Formulation

The examples of the previous section illustrate the main ideas of our verification methodology. We will now proceed with a more formal presentation, showing that these ideas apply to general classes of circuits. This section develops a mathematical abstraction of logic circuits, simulators, and the verification problem.

3.1 Notation

We adopt a notation that represents the system input, output, and state values as vectors. The history of inputs applied to a system is denoted by a sequence of vectors.

p: the number of system inputs.

B: $\{0, 1\}$, the Boolean domain.

B^n: $\{\langle x_1, \ldots, x_n \rangle | x_i \in B\}$, Boolean vectors of size n.

Φ_l: all length l sequences with elements in B^p.

Φ: $\bigcup_{0 \leq l < \infty} \Phi_l$, all finite sequences with elements in B^p.

T: $\{0, 1, X\}$, the ternary domain, partially ordered $X < 1$ and $X < 0$.

T^n: $\{\langle x_1, \ldots, x_n \rangle | x_i \in T\}$, ternary vectors of size n, partially ordered $\vec{x} \leq \vec{y}$ if $x_i \leq y_i$ for all $1 \leq i \leq n$.

11

Ψ_l: all length l sequences with elements in T^p.

Ψ: $\bigcup_{0 \le l < \infty} \Psi_l$, all finite sequences with elements in T^p, partially ordered $[\vec{a}_1, \ldots, \vec{a}_s] \le [\vec{b}_1, \ldots, \vec{b}_t]$ if $s \le t$ and $\vec{a}_{s-i} \le \vec{b}_{t-i}$ for all $0 \le i < s$.

ϵ: the empty sequence.

$\alpha \cdot \beta$:   the concatenation of sequences α and β.

3.2   Information Ordering

The partial ordering $X < 0$ and $X < 1$ orders values by their "information content." That is, X indicates an absence of information while 0 and 1 represent specific, fully-defined values. When speaking of domains ordered by information content values a and b are said to be "consistent" if either $a \le b$ or $b \le a$, and "inconsistent" otherwise. Value a is "weaker" than b if $a < b$, i.e., $a \le b$ and $a \ne b$.

The information ordering is extended to vectors and vector sequences by adopting the convention that one value is less than another if the elements of the first are consistent with those of the second, but the first contains less information. More precisely, for vectors in T^n, one vector is less than or equal to another if each element of the first is less than or equal to each element of the second. For sequences $\alpha, \beta \in \Psi$, α is less than or equal to β if the elements of α are less than or equal to the corresponding final elements of β. That is, the history given by α is consistent with the most recent history given by β, but may contain less information. The motive for this convention on the ordering of different length sequences will become clear when we study the monotonicity properties of the simulator. As a special case, sequences of Boolean vectors, $\alpha, \beta \in \Phi$, are ordered $\alpha \le \beta$ when α is a suffix of β.

Definition 1 *For partially ordered sets D_1, D_2 a monotonic function $g: D_1 \to D_2$ satisfies*

$$a \le b \Longrightarrow g(a) \le g(b)$$

for all $a, b \in D_1$. Similarly, a monotonic function with multiple arguments satisfies this property for each argument.

For any program that processes data ordered by information content, such as a logic simulator, monotonicity expresses an important property. Suppose the program is given a stimulus containing incomplete information, e.g., having some inputs equal to X. If the program obeys monotonicity, it will produce a response consistent with but possibly weaker than the response it would produce given a stronger stimulus.

12

Definition 2 *For partially ordered sets D_1 and D_2 with subsets D_1' and D_2', respectively, a monotonic extension of function $f: D_1' \to D_2'$ is a function $g: D_1 \to D_2$ such that g is monotonic and $f(a) = g(a)$ for all $a \in D_1'$. Similarly, a monotonic extension of a multiple argument function must satisfy these properties for each argument.*

As an example, the OR function can be extended monotonically from the Boolean to the ternary domain in several different ways, including the following:

$$a \vee b = \begin{cases} 1, & a = 1 \text{ or } b = 1 \\ 0, & a = 0 \text{ and } b = 0 \\ X, & \text{else.} \end{cases} \tag{1}$$

Other extensions yield X when $a = 1$ and $b = X$ or vice-versa. These more pessimistic extensions are still monotonic but would tend to cause false negative results.

3.3 System Specification

The system to be implemented has p inputs and m outputs, each of which may equal 0 or 1. Hence the system may be described as a finite automaton with input alphabet B^p and output alphabet B^m.

Definition 3 *A system specification $\mathcal{M}$ is a triple $\langle Q, Next, Out \rangle$ with*

    Q: a finite set of states,

    $Next$:   the next state function $Next: Q \times B^p \to Q$, and

    Out: the output function $Out: Q \to B^m$.

Function $Next$ must be a surjection, that is, for every $q \in Q$, there must be a $q' \in Q$ and $\vec{a} \in B^p$ such that $q = Next(q', \vec{a})$.

A system with a surjective next state function cannot have any unreachable states. This restriction is imposed for technical reasons but should not limit the class of actual systems under consideration, because there would be no reliable way to put the system in a state that cannot be reached by any sequence of state transitions. This restriction is much milder than the strong connectivity property assumed by Moore [18].

The function $Next$ is extended to input sequences to yield the state after all inputs have been applied, i.e., to a function $Next: Q \times \Phi \to Q$ defined recursively as

$$Next(q, \epsilon) \quad = \quad q$$

$$Next(q, \alpha \cdot \vec{x}) \quad = \quad Next(Next(q, \alpha), \vec{x}).$$

13

Proposition 1 *For any $k \geq 0$, The extended function $Next: Q \times \Phi_k \to Q$ is a surjection.*

Proof: By induction on k and the surjectivity of *Next* when applied to single inputs.

□

The function $FinalOut: Q \times \Phi \to B^m$ is defined to yield the final output after a sequence of inputs has been applied, i.e.,

$$FinalOut(q, \alpha) = Out(Next(q, \alpha)).$$

Proposition 2

$$Next(q, \alpha \cdot \beta) = Next(Next(q, \alpha), \beta)$$

and

$$FinalOut(q, \alpha \cdot \beta) = FinalOut(Next(q, \alpha), \beta)$$

Proof: By induction on the length of β.

□

3.4 Circuit Model

A circuit is also a finite automaton but with states encoded by s Boolean variables.

Definition 4 *A circuit C is a triple $\langle s, next, out \rangle$ with*

   s: the number of state variables,

   next: the next state function $next: B^s \times B^p \to B^s$, and

   out: the output function $out: B^s \to B^m$.

This definition of a circuit differs significantly from that assumed by other verifiers. Whereas structural verifiers describe a circuit as a set of interconnected elements, our abstraction ignores the physical structure altogether. Instead, it views a circuit at the level seen by the user of a logic simulator. That is, the program, rather than the user, is responsible for determining the behavior of the circuit given its structure. This level of detail suffices for the paper, because our goal in verification is only to ensure that the user uncovers any design errors that can be detected by the simulator.

Definition 5 *Circuit C implements* specification *M when there exists a relation $\mathcal{E} \subseteq Q \times B^s$ (for "encodes") satisfying:*

1. *For every $q \in Q$ there exists a $\vec{z} \in B^s$ for which $q \,\mathcal{E}\, \vec{z}$.*

2. *For any $q \in Q$ and $\vec{z} \in B^s$*

$$q \,\mathcal{E}\, \vec{z} \implies Out(q) = out(\vec{z}),$$

3. *For any $q \in Q$, $\vec{z} \in B^s$, and $\vec{x} \in B^p$,*

$$q \,\mathcal{E}\, \vec{z} \implies Next(q, \vec{x}) \,\mathcal{E}\, next(\vec{z}, \vec{x}).$$

By this definition, the circuit automaton must "cover" the input–output behavior of the specification. That is, for any initial state of the specification, there must be an initial circuit state, such that the two automata would yield identical outputs for any subsequent input sequence. However, there may be circuit states that do not correspond to any specification states, such as those involving invalid combinations of state variables. Furthermore, neither automata need be reduced—several circuit states may correspond to a single specification state and vice-versa.

3.5   Simulator

For a circuit $C = \langle s, next, out \rangle$, the simulator maintains state variables $\vec{z} \in T^s$, and computes the behavior according to functions $next \colon T^s \times T^p \to T^s$ and $out \colon T^s \to T^m$, which are arbitrary monotonic extensions of the corresponding circuit functions. The simulator implements five commands, although for black-box simulation only the first three are allowed.

ERASE: causes the simulator to set z_i to X for $1 \le i \le s$.

CYCLE($\vec{x}$):   causes the simulator to set $\vec{z}$ to $next(\vec{z}, \vec{x})$.

OUTPUT: causes the simulator to print $out(\vec{z})$.

SET(i, b):   causes the simulator to set z_i to b for $b \in T$.

OBSERVE(i):   causes the simulator to print z_i.

15

3.6 Simulation Experiment

Definition 6 *A* simulation experiment *consists of a sequence of simulation commands beginning with* ERASE, *as well as a procedure by which the user decides whether the outcome is acceptable. The decision procedure can depend only on the values produced by* OUTPUT *and* OBSERVE *commands.*

Definition 7 *An experiment is* effective *for specification* $\mathcal{M}$ *when the simulation of a circuit* C *can be judged acceptable only if* C *implements* $\mathcal{M}$.

This condition requires that the verification method cannot produce any false positive responses.

Definition 8 *An experiment is* nontrivial *when the simulator produces an acceptable outcome for some circuit* $\langle s, next, out \rangle$ *and monotonic extensions of next and out.*

This condition is imposed to eliminate the otherwise effective test of rejecting all circuits.

4 Black-Box Simulation

With black-box simulation, the user is limited to the simulation commands ERASE, CYCLE, and OUTPUT. No direct observation or modification of the simulator state $\vec{z}$ is permitted. This section identifies the class of systems that can be verified by black-box simulation.

Define the function $SimState: \Psi \to T^{s}$ as the state of the simulator after giving an ERASE command followed by a series of CYCLE commands. More precisely

$$SimState(\epsilon) \quad = \quad X^{s}$$

$$SimState(\alpha \cdot \vec{x}) \quad = \quad next(SimState(\alpha), \vec{x})$$

where X^{s} denotes a vector of size s with each element equal to X.

Similarly, define the function $SimOut: \Psi \to T^{m}$ as the result that would be printed by an OUTPUT command following the simulation of some input sequence, i.e.,

$$SimOut(\alpha) = out(SimState(\alpha)).$$

A black-box simulation experiment can be viewed as a decision procedure that either accepts or rejects a circuit based on the values of $SimOut(\alpha)$ for any finite number of sequences $\alpha \in \Psi$.

4.1 Definite Systems

Definition 9 *For any $k \geq 0$, specification $\mathcal{M}$ is k-definite if*

$$FinalOut(q_1, \alpha) = FinalOut(q_2, \alpha)$$

for any $\alpha \in \Phi_k$ and any $q_1, q_2 \in Q$,

Definition 10 *A specification is* definite *if it is k-definite for any k. Otherwise it is* indefinite.

This class of sequential systems was first identified by Kleene [13]. Since that time, various definitions have appeared, viewing sequential systems either as recognizers [19] or transducers [14]. Our definition most closely matches that of Kohavi [14]. However, he defines a k-definite system as one for which any input sequence of length k places the system in a unique state, whereas we only require the sequence to cause a unique output. If the specification automaton is in reduced form, the two definitions are equivalent.

Proposition 3 *A k-definite specification is also l-definite for any $l \geq k$.*

Proof: Any sequence in Φ_l is of the form $\beta \cdot \alpha$ where $\beta \in \Phi_{l-k}$ and $\alpha \in \Phi_k$. For any states $q_1, q_2 \in Q$, let $q_1' = Next(q_1, \beta)$ and $q_2' = Next(q_2, \beta)$. By Proposition 2 and the fact that the specification is k-definite:

$$FinalOut(q_1, \beta \cdot \alpha) = FinalOut(q_1', \alpha) = FinalOut(q_2', \alpha) = FinalOut(q_2, \beta \cdot \alpha)$$

$\square$

4.2 Monotonicity Properties

In this section we will prove several properties of the logic simulator that follow from the monotonicity of the simulator functions *next* and *out*. As shall be seen, monotonicity provides the primary mechanism by which one can guarantee properties of a circuit knowing only its response during simulation.

Lemma 1 *The functions SimState and SimOut are monotonic.*

Proof: We will prove by induction on the length of α that for any $\alpha, \beta \in \Psi$ for which $\alpha \leq \beta$, we have $SimState(\alpha) \leq SimState(\beta)$. First, if $\alpha = \epsilon$, then $SimState(\alpha) = X$'

and this vector is less than or equal to any other state vector. Otherwise, if α has nonzero length and $\alpha \leq \beta$, then α must be of the form $\alpha' \cdot \vec{a}$ and β must be of the form $\beta' \cdot \vec{b}$ where $\alpha' \leq \beta'$ and $\vec{a} \leq \vec{b}$. Assuming, by induction, that $SimState(\alpha') \leq SimState(\beta')$ and given that $next$ is monotonic, we get

$$SimState(\alpha) = next(SimState(\alpha'), \vec{a}) \leq next(SimState(\beta'), \vec{b}) = SimState(\beta).$$

The monotonicity of $SimOut$ follows from the fact that both out and $SimState$ are monotonic, because a composition of monotonic functions is also monotonic.

□

The monotonicity of $SimState$ and $SimOut$ show how the ERASE command and the three-valued modeling enhances the power of the simulator. If an ERASE command followed by a sequence of CYCLE commands causes the simulator to produce a 0 or 1 on some output or state variable, then this sequence of inputs must also cause the circuit to produce the same output or state regardless of the initial state. As a special case of this lemma, if sequence α' is a suffix of α, then $SimOut(\alpha') \leq SimOut(\alpha)$, i.e., the simulation of the shorter sequence yields an output consistent with, but possibly weaker than, the output produced for the longer one.

Lemma 2 *For any $k \geq 0$, if $SimOut(\alpha) = FinalOut(q, \alpha)$ for all $\alpha \in \Phi_k$ and all $q \in Q$, then the simulated circuit C implements specification M.*

Proof: For any $\alpha \in \Phi_k$ let

$$Q_\alpha = \{Next(q', \alpha) | q' \in Q\}$$

and

$$Z_\alpha = \{\vec{z} \in B^s | SimState(\alpha) \leq \vec{z}\}$$

That is Q_α denotes the set of possible states for the specification automaton following input sequence α, while Z_α denotes the set of possible circuit states consistent with the state of the simulator after simulating the sequence α following an ERASE command. By the monotonicity of out, we must have that for $\vec{z} \in Z_\alpha$

$$SimOut(\alpha) = out(SimState(\alpha)) \leq out(\vec{z}).$$

For those sequences α that occur in the condition of the lemma, $SimOut(\alpha)$ is maximal, i.e., $SimOut(\alpha) \in B^m$, in which case we can conclude that $out(\vec{z}) = SimOut(\alpha)$ for all $\vec{z} \in Z_\alpha$.

Define $\mathcal{E}$ as

$$\mathcal{E} = \bigcup_{\alpha \in \Phi_k} \{(q, \vec{z}) | q \in Q_\alpha, \vec{z} \in Z_\alpha\}.$$

We must show that $\mathcal{E}$ satisfies the three properties of Definition 5.

18

First, given that *Next* is surjective for input sequences of length k (Proposition 1), every state q must be in set Q_α for some $\alpha \in \Phi_k$. The set Z_α cannot be empty, and hence $q \mathcal{E} \vec{z}$ for some $\vec{z}$.

Second, if $q \mathcal{E} \vec{z}$, we must have $q = Next(q', \alpha)$ and $SimState(\alpha) \leq \vec{z}$ for some $q' \in Q$ and some $\alpha \in \Phi_k$. From the condition of the lemma it follows that

$$Out(q) = FinalOut(q', \alpha) = SimOut(\alpha) = out(\vec{z}).$$

Finally, suppose $q \mathcal{E} \vec{z}$, i.e., for some $\alpha = [\vec{a}_1, \ldots, \vec{a}_k]$ we have $q \in Q_\alpha$ and $SimState(\alpha) \leq \vec{z}$. Consider any $\vec{x} \in B^p$, and let $\gamma = [\vec{a}_1, \ldots, \vec{a}_k, \vec{x}]$ and $\beta = [\vec{a}_2, \ldots, \vec{a}_k, \vec{x}]$. By definition, $Next(q, \vec{x}) \in Q_\beta$. Since $\beta \leq \gamma$ (β is a suffix of γ), and both $SimState$ and $next$ are monotonic

$$SimState(\beta) \leq SimState(\gamma) = next(SimState(\alpha), \vec{x}) \leq next(\vec{z}, \vec{x})$$

Therefore $next(\vec{z}, \vec{x}) \in Z_\beta$ by the definition of Z_β and hence $Next(q, \vec{x}) \mathcal{E} next(\vec{z}, \vec{x})$.
□

This lemma provides the key to proving that a simulator can verify that a circuit implements a k-definite specification by simulating it for all input sequences of length k.

Lemma 3 *If circuit C implements specification M then for all $\alpha \in \Psi$ and all $\beta \in \Phi$ such that $\alpha \leq \beta$:*

$$SimOut(\alpha) \leq FinalOut(q, \beta)$$

for all $q \in Q$.

Proof: We will prove by induction on the length of β that for some $\vec{z} \in B^s$ such that $Next(q, \beta) \mathcal{E} \vec{z}$, we have $SimState(\beta) \leq \vec{z}$. Given this, we can infer by the monotonicity of $SimOut$ and out, and by condition 2 of Definition 5 that

$$SimOut(\alpha) \leq SimOut(\beta) = out(SimState(\beta)) \leq out(\vec{z}) = FinalOut(q, \beta).$$

To prove the induction hypothesis, for $\beta = \epsilon$, we have that $SimState(\beta) = X^s$ and hence the hypothesis holds trivially. Now suppose that β is of the form $\beta = \beta' \cdot \vec{b}$, that $Next(q, \beta') \mathcal{E} \vec{z}'$, and $SimState(\beta') \leq \vec{z}'$. By definition

$$Next(q, \beta) = Next(Next(q, \beta'), \vec{b})$$

and therefore by condition 3 of Definition 5

$$Next(q, \beta) \mathcal{E} next(\vec{z}', \vec{b}).$$

By the monotonicity of *next*

$$SimState(\beta) = next(SimState(\beta'), \vec{b}) \leq next(\vec{z}', \vec{b})$$

Hence, if we let $\vec{z} = next(\vec{z}', \vec{b})$ the induction hypothesis will hold.

□

This lemma implies that if the specification has states q_1 and q_2 for which $FinalOut(q_1, \alpha) \neq FinalOut(q_2, \alpha)$, for some sequence α, then some element of $SimOut(\alpha)$ must equal X even for a correctly designed circuit. This property is used in designing a booby trap for an indefinite system.

4.3 Expressive Power

We are now ready to prove a main result of this paper, characterizing the capabilities and limitations of black-box simulation.

Theorem 1 *There exists an effective, nontrivial, black-box simulation experiment for specification M if and only if M is definite.*

Proof: First, suppose M is k-definite for some value k. Consider the simulation experiment consisting of executing the sequence of commands required to compute $SimOut(\alpha)$ for each $\alpha \in \Phi_k$, and accepting the circuit if $SimOut(\alpha) = FinalOut(q, \alpha)$ for all α and any choice of $q \in Q$ (in a k-definite specification, the choice of initial state makes no difference.) Lemma 2 shows that this experiment is effective.

Furthermore, the circuit illustrated in Figure 4, consisting of a p-bit wide, k-bit long shift register to store the most recent k inputs plus logic to compute the circuit outputs can pass this experiment. A similar structure was proposed by Kleene [13] to implement an arbitrary definite system. More precisely, let $s = p \cdot k$ and define *next* as

$$next_i(\vec{z}, \vec{x}) = \begin{cases} x_{i-p \cdot (k-1)}, & p \cdot (k-1) < i \leq p \cdot k \\ z_{i+p}, & 1 \leq i \leq p \cdot (k-1) \end{cases} \tag{2}$$

Partition $\vec{z}$ into a sequence of vectors $[\vec{z}_1, \ldots, \vec{z}_k]$ where $\vec{z}_i = \langle z_{(i-1) \cdot p+1}, \ldots, z_{i \cdot p} \rangle$ and define *out* as

$$out(\vec{z}) = FinalOut(q, [\vec{z}_1, \ldots, \vec{z}_k])$$

for any choice of $q \in Q$. These functions are extended monotonically by defining $next(\vec{z}, \vec{x})$ according to Equation 2 and letting *out* be any monotonic extension. It can be seen that for any sequence $\alpha = [\vec{a}_1, \ldots, \vec{a}_k] \in \Phi_k$, we have $SimState_{(i-1) \cdot p+j}(\alpha) = a_{i,j}$, the jth element of vector $\vec{a}_i$, and hence $SimOut(\alpha) = FinalOut(q, \alpha)$. Therefore, the simulation experiment is nontrivial.

20

Figure 4: Universal Implementation of a k-Definite System

Next, suppose M is not k-definite for any value of k. We use an adversary argument to show that no nontrivial, black-box experiment for this specification can also be effective. Assume that there is some nontrivial simulation experiment in which fewer than k CYCLE commands occur without an intervening ERASE command, for some value k. Since the experiment is nontrivial, there must be a circuit $C = \langle s, next, out \rangle$ and monotonic extensions of $next$ and out that produce an acceptable result for the experiment. We will construct a circuit $C' = \langle s', next', out' \rangle$ that does not implement M. However, we can define monotonic extensions of $next'$ and out' such that by defining $SimOut'$ in a manner analogous to the definition of $SimOut$, we have $SimOut'(\alpha) = SimOut(\alpha)$ for any sequence $\alpha \in \Psi_l$, for which $l < k$. The simulation experiment cannot possibly distinguish C from C' and hence is not effective. This argument holds for any value of k, showing that there is no finite upper bound on the length of a simulation experiment that can distinguish a correct circuit from a defective one when the specified system is indefinite.

Circuit C' is constructed as illustrated in Figure 5 by taking circuit C and adding extra logic to implement a "booby trap", i.e., logic that will not affect the output value until a specific input sequence of length k occurs. Designing such a booby trap is no easy task, because any state variables used by the trap will be set to X whenever the user gives an ERASE command. For an improper design this could cause C' to produce an X on its output under conditions when C would not. Similarly, the user might attempt to expose any traps by presenting inputs with some elements equal to X.

21

Figure 5: Circuit with Booby Trap

For the design of C', let $\beta = [\vec{b}_1, \ldots, \vec{b}_k] \in \Phi_k$ be some sequence such that $\mathcal{M}$ has states $q_1, q_2 \in Q$ for which $FinalOut(q_1, \beta) \neq FinalOut(q_2, \beta)$. Such a sequence must exist or else $\mathcal{M}$ would be k-definite. Assume for simplicity, that output m differs for these two cases, relabeling the outputs if required. The trap consists of a shift register, where each shift element sets its output to its input value when the circuit input matches the corresponding element of β and clears the output to 0 otherwise. Consequently, input sequence β will cause a 1 to propagate through the shift register, forcing output m to 1 when it reaches the end. Any input sequence $\alpha \not\leq \beta$ of length less than or equal to k will cause the shift register to produce 0, leaving the circuit output unchanged. The behavior of the circuit for input sequences $\alpha < \beta$ will depend on the initial state, and hence under such conditions $SimOut'_m(\alpha) = X$. However, it can be shown using Lemma 3 that under these conditions $SimOut_m(\alpha) = X$ as well.

The detailed design of circuit C' is rather involved and hence is given in Appendix A rather than here. It is also shown that for any $l \leq k$, any $\alpha \in \Psi_l$, and any i such that $1 \leq i \leq m$:

$$SimOut'_i(\alpha) = \begin{cases} 1, & i = m \text{ and } \alpha = \beta \\ SimOut_i(\alpha), & \text{else} \end{cases}$$

Hence C' cannot be distinguished from C for any input sequence of length less than k. On the other hand, since $FinalOut_m(q_1, \beta) \neq FinalOut_m(q_2, \beta)$, we must have either $SimOut'_m(\beta) \not\leq FinalOut_m(q_1, \beta)$ or $SimOut'_m(\beta) \not\leq FinalOut_m(q_2, \beta)$ and hence by Lemma 3, C' cannot implement $\mathcal{M}$.

□

22

5 State Transition Simulation

To verify circuits implementing a more general class than definite systems, the results of the previous section imply that a capability beyond black-box simulation is required. At the opposite extreme, if the user were to completely specify the relation $\mathcal{E}$, we could check that it satisfies the conditions of Definition 5 by exhaustively simulating all states and transitions. This approach would work for any class of systems and circuits. In practice, however, it cannot be applied to circuits of significant size, because the complexity of completely specifying and checking the relation $\mathcal{E}$ would be overwhelming. Exhaustive simulation, however, provides a basis for developing other simulation methods that overcome the deficiencies of black-box simulation. It shows that any circuit can be verified if we introduce sufficient detail about the circuit structure into the verification method.

5.1 The Assertion Method

We would prefer to introduce as little information as possible about the circuit structure into the verification. Toward this goal we will develop a notation similar to the Floyd-Hoare assertion method of program verification, along with an associated simulation methodology for testing assertions.

Definition 11 *A* circuit assertion *is a set of four predicates:*

$P(\vec{z})$:   *a precondition on the state variables,* $P: B^s \rightarrow \{true, false\}$.

$I(\vec{x})$:   *a condition on the input variables,* $I: B^p \rightarrow \{true, false\}$.

$N(\vec{z})$:   *a postcondition on the state variables,* $N: B^s \rightarrow \{true, false\}$.

$O(\vec{y})$:   *a postcondition on the output variables,* $O: B^m \rightarrow \{true, false\}$.

A circuit assertion is denoted by an equation of the form the form $P\{\ I\ \}N \wedge O$.

Definition 12 *A circuit satisfies assertion* $P\{\ I\ \}N \wedge O$ *if* $N[next(\vec{x}, \vec{z})]$ *and* $O[out(next(\vec{x}, \vec{z}))]$ *hold for all* $\vec{z} \in B^s$ *and* $\vec{x} \in B^p$ *such that* $P(\vec{z})$ *and* $I(\vec{x})$ *hold.*

Given a set of circuit assertions, verifying a circuit requires two proofs—that any circuit satisfying the set of assertions must implement the specification, and that the circuit under consideration satisfies these assertions. Proving the adequacy of a set of assertions involves showing that they cover every transition in the specification automaton. At the present

stage of this research, the set of assertions and a proof of their adequacy must be generated manually. Although this places additional burden on the user, experience has shown that far less manual effort is required than with structural verifiers.

5.2   Testing Assertions by Simulation

Once a set of assertions has been devised, a simulator can verify that a particular circuit satisfies them. First, we must structure the assertions in a particular way.

Definition 13 *A nonvoid predicate* $P: B^n \rightarrow \{true, false\}$ *is* convex *if whenever* $P(\vec{a})$ *and* $P(\vec{b})$ *hold for vectors* $\vec{a}$ *and* $\vec{b}$, *then* $P(\vec{c})$ *holds for any vector* $\vec{c}$ *for which* $c_i \in \{a_i, b_i\}$ *for* $1 \leq i \leq n$.

A convex predicate can be expressed by a formula of the form $P(\vec{u}) = L_1 \wedge L_2 \wedge \ldots \wedge L_k$ where each L_i is a *literal* of the form $u_j = 0$ or $u_j = 1$. In general, any assertion can be rewritten as a set of assertions containing only convex predicates.

Definition 14 *For convex predicate* $P(\vec{u})$ *defined over elements of* B^n, *the vector* $\vec{u}_P \in T^n$ *is defined as*

$$[\vec{u}_P]_i = \begin{cases} 1, & P(\vec{a}) \Rightarrow a_i = 1 \\ 0, & P(\vec{a}) \Rightarrow a_i = 0 \\ X, & \text{else} \end{cases}$$

For a predicate expressed as a conjunction of literals, the corresponding vector is obtained by setting each element appearing in a literal to its specified value, and all other elements to X. This vector is analogous to the cubical representation of a product term in a Boolean expression [20].

Proposition 4 *For convex predicate* $P(\vec{u})$ *and vector* $\vec{a} \in B^n$, $P(\vec{a})$ *if and only if* $\vec{u}_P \leq \vec{a}$.

Proof: First, assume that $P(\vec{a})$ holds. For the vectors to be ordered $\vec{u}_P \not\leq \vec{a}$, there must be at least one element i such that $[\vec{u}_P]_i = \neg a_i$. This, however, would violate Definition 14, and hence $\vec{u}_P \leq \vec{a}$.

On the other hand, assume $\vec{u}_P \leq \vec{a}$ for some $a \in B^n$. Let $\vec{b} \in B^n$ be a vector such that $P(\vec{b})$ holds. The following procedure constructs a sequence of vectors $\vec{b}^0, \vec{b}^1, \ldots, \vec{b}^n$, such that $\vec{b} = \vec{b}^0$, $\vec{a} = \vec{b}^n$, and $P(\vec{b}^i)$ holds for $0 \leq i \leq n$. From this we can conclude that $P(\vec{a})$ holds. For $1 \leq i \leq n$ each element j of vector $\vec{b}^i$ defined as

$$b_j^i = \begin{cases} b_j^{i-1}, & j \neq i \\ a_j, & j = i \end{cases}$$

24

The proof that each vector $\vec{b}^i$ satisfies $P(\vec{b}^i)$ proceeds by induction on i. Since $\vec{b}^0 = \vec{b}$ it clearly holds for the basis case. Assuming $\vec{b}^{i-1}$ satisfies $P(\vec{b}^{i-1})$, observe that $\vec{b}^{i-1} \neq \vec{b}^i$ only if $a_i = \neg b_i^{i-1}$ and $[\vec{u}_P]_i = X$. There must be some vector $\vec{d} \in B^n$ such that $P(\vec{d})$ and $d_i = \neg b_i^{i-1}$, or else $[\vec{u}_P]_i$ would equal b_i^{i-1}. By convexity, $P(\vec{b}^i)$ must hold, because each element of $\vec{b}^i$ equals an element of $\vec{d}$ or an element of $\vec{b}^{i-1}$.

$\square$

This result shows that a convex predicate can be represented by a single cube.

Theorem 2 *For convex predicates P, I, N, and O if $\vec{z_N} \leq next(\vec{z}_P, \vec{x}_I)$ and $\vec{y_O} \leq out(next(\vec{z}_P, \vec{x}_I))$ for any monotonic extension of circuit functions next and out, then the circuit satisfies assertion $P\{\ I\ \}N \wedge O$.*

Proof: Suppose $\vec{z_N} \leq next(\vec{z}_P, \vec{x}_I)$. By Proposition 4, any vector $\vec{z} \in B^s$ for which $P(\vec{z})$ holds must satisfy $\vec{z_P} \leq \vec{z}$. Similarly, any vector $\vec{x} \in B^p$ for which $I(\vec{x})$ holds must satisfy $\vec{x_I} \leq \vec{x}$. Therefore, the monotonicity of *next* implies that

$$\vec{z_N} \leq next(\vec{z}_P, \vec{x}_I) \leq next(\vec{z}, \vec{x})$$

and $N[next(\vec{x}, \vec{z})]$ holds by Proposition 4.

Similarly, when $\vec{y_O} \leq out(next(\vec{z}_P, \vec{x}_I))$ the monotonicity of *next* and *out* imply that

$$\vec{y_O} \leq out(next(\vec{z}_P, \vec{x}_I)) \leq out(next(\vec{z}, \vec{x}))$$

and hence $O[out(next(\vec{x}, \vec{z}))]$ holds.

$\square$

This theorem indicates a straightforward procedure to test that a circuit satisfies an assertion with convex predicates. Following an ERASE command, use SET commands to set all state variables for which $[\vec{z}_P]_i \neq X$ to the appropriate values. Then give the command CYCLE$(\vec{x}_I)$ to simulate the prescribed action. Finally, use OBSERVE commands to check that $z_i = [\vec{z}_N]_i$ for all i such that $[\vec{z}_N]_i \neq X$, and an OUTPUT command to check that $\vec{y_O} \leq out(\vec{z})$.

6 Performance Considerations

Up to this point, we have considered only whether verifying a circuit by simulation was at all possible. The resulting verification methods were not at all efficient. For example, brute force application of black-box simulation to verify a k-definite system with m inputs requires simulating 2^{km} patterns. Clearly, this is practical only for small values of k and m. In general, the circuit verification problem is NP-hard as measured in the size of the circuit and the specification. However, several techniques reduce the complexity to manageable levels for a large class of circuits.

6.1 Input Weakening

The logic value X can be used to indicate a "don't care" (or more properly "shouldn't care") condition when the circuit behavior being tested should not depend on that particular input. This allows us to simulate the effects of a number of Boolean sequences with a single ternary sequence, leading at times to a dramatic reduction in the simulation complexity. This technique is called "input weakening", because it involves reducing the information content of the simulation sequences. Monotonicity guarantees that if the resulting response on some output is 0 or 1, then all stronger sequences would give the same response.

For example, consider a k-bit long, 1-bit wide shift register. Brute force, black-box simulation requires simulating 2^k patterns of the form $[a_1, a_2, \ldots, a_k]$, each time checking that the final output equals a_1. Since the output of the shift register should depend only on the first value in the sequence, we can set the input to X for the remainder of the simulation. This reduces the number of simulation sequences to two: $[1, X, \ldots, X]$ and $[0, X, \ldots, X]$, without compromising the rigor of the simulation. Generalizing this to a shift register of width m, a total of $2m$ sequences, each of length k, suffices, consisting of a pair to test each bit of the data word. Compared to the *ad hoc* methods most designers use to validate shift registers (e.g., simulate a randomly chosen input sequence), the proposed method provides better results at a comparable cost.

To develop this idea formally, we define a covering set as a set of ternary sequences that include all possible Boolean sequences of a given length.

Definition 15 *A set $A \subseteq \Psi$ is a* covering set *for Φ_k if to every $\beta \in \Phi_k$ there corresponds some $\alpha \in A$ such that $\alpha \leq \beta$.*

Theorem 3 *For a covering set A of Φ_k, if $SimOut(\alpha) = FinalOut(q, \beta)$ for all $\alpha \in A$, all $\beta \in \Phi_k$ such that $\alpha \leq \beta$, and all $q \in Q$, then the simulated circuit C implements specification $\mathcal{M}$.*

Proof: By the monotonicity of $SimOut$, if $\alpha \leq \beta$, then $SimOut(\alpha) \leq SimOut(\beta)$. However, the assumption that $SimOut(\alpha) = FinalOut(q, \beta)$ implies that $SimOut(\alpha) \in B^m$, and hence $SimOut(\beta) = FinalOut(q, \beta)$. Thus, the conditions required by Lemma 2 hold, and C implements $\mathcal{M}$.

$\square$

Input weakening can also be applied in transition simulation. In fact the simulation sequences arising from the assertion method already utilize this technique. If convex predicate P places no conditions on element i of a vector, then corresponding variable is set to X in the simulation.

6.2 Symbolic Simulation

At times we cannot avoid the complexity caused by the large number of possible input combinations that might be applied to a circuit, all of which might be relevant to the values of the outputs. For these cases, we propose *symbolic simulation* to reduce the number of patterns simulated. A symbolic simulator [5] resembles a conventional logic simulator, except that the input sequences can contain Boolean variables in addition to the constants 1 and 0. During simulation the values of the circuit state and output are Boolean functions of the variables occurring in the input sequence. A symbolic simulator represents and manipulates these functions explicitly. The worst-case behavior of such a program gives no better performance than exhaustive simulation by a conventional simulator. However, good Boolean manipulation algorithms often lead to far better results. To implement the verification methodologies described in this paper, a symbolic simulator must be able to manipulate functions over the three-valued domain $\{0, 1, X\}$. The symbolic simulator MOSSYM [5] solves this problem by representing every circuit variable by a pair of Boolean functions, generalizing the encoding of three possible values by two bits.

A symbolic simulator can verify an m-input, k-definite system by simulating a single sequence of length k, with each input pattern consisting of m Boolean variables to represent all possible input values. The resulting output functions will be symbolic representations of the circuit outputs for every possible input sequence of length k. These can then be tested for equivalence with functions generated from the system specification. As an example of verification by symbolic simulation, the above-mentioned shift register would be verified by simulating the sequence of variables $[a_1, a_2, \ldots, a_k]$ and testing the final output for equivalence with the function a_1. Efficient symbolic manipulation will exploit the fact that the variables shifting through the register do not interact. Hence a symbolic simulator can automatically take advantage of the same properties that allow input weakening. Symbolic simulation can also handle cases for which input weakening does not apply. For example, MOSSYM was able to verify a 16-bit nMOS adder using less than 10 minutes of CPU time on a Digital Equipment Corporation VAX-11/780. In contrast, its more traditional counterpart MOSSIM II [4] would require an estimated 648 years using exhaustive black-box simulation.

The capabilities of symbolic simulator can also be exploited in verifying indefinite systems. Rather than testing a large number of assertions with predicates containing terms of the form $v = 0$ or $v = 1$, the program would test a smaller set of assertions containing predicates of the form $v = a$ where a is a symbolic variable. That is, a symbolic simulator can test assertions having the form of universally quantified formulas rather than single propositions. Furthermore, it does not require the predicates to be convex.

Although a symbolic simulator gives the user a far more abstract view of circuit behavior, it has no fundamental power beyond that of an ordinary logic simulator. Any information that symbolic simulation provides could also be obtained by exhaustively simulating the set of patterns generated by enumerating all combinations of 0 and 1 for the Boolean

Figure 6: Static RAM Circuit

variables.

7 Memory Verification Example

To demonstrate the methodology on a more significant task for hardware verification, consider the static random-access memory (RAM) illustrated in Figure 6. The assertions required to verify this circuit will be presented in a series of steps, each introducing new notation and discussing the reasoning behind it. Despite these extensions to the notation, the underlying principle remains that of state transition simulation.

This circuit holds $n = 2^m$ bits, where each memory cell i, such that $0 \leq i < n$, consists of a feedback path containing electrical nodes B_i and $\overline{B}_i$ along with a pair of access transistors [10]. As a shorthand, the predicate $Store(i, v)$ expresses the fact that value $v \in \{0, 1\}$ is stored in memory cell i:

$$Store(i, v) \quad \equiv \quad B_i = v \ \wedge \ \overline{B}_i = \neg v.$$

The input lines $ADDR_j$, for $0 \leq j < m$, select a particular memory cell. When WRITE $= 1$, the value of DIN is written into the selected memory cell. As shorthand, this operation is expressed by a predicate

$$Write(i, v) \quad \equiv \quad \text{WRITE} = 1 \ \wedge \ \text{DIN} = v \ \wedge \ \forall (0 \leq k < m)[\text{ADDR}_k = i_k]$$

where i_k indicates the kth bit in the binary representation of i. When WRITE $= 0$, the value stored in the selected memory cell is produced as output on DOUT. This operation

28

is expressed by the predicate

$$Read(i) \quad \equiv \quad \text{WRITE} = 0 \ \wedge \ \forall (0 \le k < m) \lceil \text{ADDR}_k = i_k \rceil.$$

Although few additional details of the circuit design are needed for verification, correct circuit operation depends on the fact that the control lines WORD_r equal 0 when the circuit is quiescent, for $0 \le r < \sqrt{n}$.[1] Without this property, the access transistors for more than one cell in a column could be turned on, causing undesirable interactions. This fact is formulated as a *system invariant*

$$Inv \quad \equiv \quad \forall (0 \le r < \sqrt{n})[\text{WORD}_r = 0].$$

The invariance of this condition is expressed by a single assertion:

$$true \ \big\{ \ true \ \big\} \ Inv$$

That is, following any memory operation, the word lines will return to a quiescent condition. Once the assertion has been established, the invariant Inv can be assumed as a precondition in all other assertions. Most circuits require some form of system invariant expressing conditions about the control logic that can be assumed true at the beginning of every input cycle. Devising the invariant requires a combination of analysis and experimentation. An insufficient system invariant will become immediately apparent during subsequent simulations, because output or state variables that should have Boolean values will equal X.

The remaining assertions simply express the operation of a memory. First, for all $v \in \{0, 1\}$ and for all i such that $0 \le i < n$, an assertion states that writing v into location i must cause v to be stored in cell i:

$$Inv \ \big\{ \ Write(i, v) \ \big\} \ Store(i, v).$$

Second, for all $v \in \{0, 1\}$ and for all i and j such that $0 \le i, j < n$ and $i \ne j$, an assertion states that writing into location j does not affect the value in cell i:

$$Inv \ \wedge \ Store(i, v) \ \big\{ \ Write(j, X) \ \big\} \ Store(i, v). \tag{3}$$

Third, for all $v \in \{0, 1\}$ and for all i such that $0 \le i < n$, an assertion states that reading location i causes its value to appear on the output:

$$Inv \ \wedge \ Store(i, v) \ \big\{ \ Read(i) \ \big\} \ \text{DOUT} = v.$$

Finally, for all $v \in \{0, 1\}$ and for all i such that $0 \le i < n$, an assertion states that reading a value from any location should have no effect on the value stored in location i:

$$Inv \ \wedge \ Store(i, v) \ \big\{ \ \text{WRITE} = 0 \ \big\} \ Store(i, v). \tag{4}$$

---

[1] A memory circuit is generally configured as a square array of memory cells with $m/2$ of the address bits selecting a row and the remaining selecting a column. Hence there are $\sqrt{n}$ rows.

The above equations represent a total of $2n^2 + 6n + 1$ assertions. The number can be further reduced by exploiting input weakening for the cases covered by Equation 3. That is, for an address i with bit representation $\langle i_0, \ldots, i_{m-1} \rangle$, all addresses j such that $j \neq i$ are covered by vectors of the form $\langle X, \ldots, X, \neg i_k, X, \ldots, X \rangle$ for $0 \leq k < m$. Thus, Equation 3 can be replaced by the following set of assertions for $v \in \{0, 1\}$, $0 \leq i < n$, and $0 \leq k < m$:

$$Inv \wedge Store(i, v) \; \big\{ \; \text{WRITE} = 1 \wedge \text{ADDR}_k = \neg i_k \; \big\} \; Store(i, v). \tag{5}$$

This reduces the total number of assertions to $2n \log n + 6n + 1$. In practice, many memory circuits would yield false negative responses for some of the assertions of Equations 4 and 5. The simulation of an assertion that causes the word line of a memory cell to be set to X would most likely corrupt the value stored in the cell. With more care, however, a set of assertions can be devised that avoid this problem while maintaining the $O(n \log n)$ bound on the total number of patterns to be simulated.

Considering that even a minimal validation of a memory circuit requires simulating $\Omega(n)$ patterns (e.g., read and write every memory location), simulating $O(n \log n)$ patterns seems a very reasonable price to pay for rigorous verification. The efficiency of this verification results from an extreme form of input (and state variable) weakening. The verification isolates each memory location, proving that it can be written and read properly, and that operations on other memory locations do not corrupt its stored value. During each test, those memory locations not under consideration are set to X. Should the circuit contain an undesirable pattern sensitivity, at least one of the tests will fail with an output or state variable equal to X that should equal 0 or 1.

Devising a simulation sequence to verify this memory requires paying a significant amount of attention to the details of the circuit design, especially when trying to minimize the number of assertions. In contrast to black-box simulation, the resulting simulation sequence is highly circuit dependent. Compare this effort, however, to that required by other verifiers. For structural verification the user would be required to specify the operation of all aspects of the circuit including the address decoders, bit lines, sense amps, and control logic. These specifications would require a circuit model that can express such effects as bidirectional transistor behavior, ratioed circuits, and precharged logic. By comparison, behavioral verification seems quite straightforward.

8   Discussion

This paper has outlined a method for applying three-valued logic simulation to the task of hardware verification. Complex circuits can be rigorously verified given only information about the desired input-output behavior, and possibly some information about the circuit state variables.

Multi-valued modeling provides the fundamental mechanism by which a circuit can be

verified knowing little about its internal structure. The requirements placed on the simulator to support verification are fairly mild. Most contemporary logic simulators provide a value X to avoid the need to find an initial Boolean state of the circuit that does not cause oscillations [12]. Although an explicit ERASE command may not be provided, it can easily be implemented, or the same result can be obtained by simply restarting the program. The monotonicity requirement simply expresses the desirable property that in the presence of X values, the simulator should not set an output or state variable to 0 or 1, when this would not have occurred had some of the X's been 0 or 1 instead. Any reasonable implementation satisfies this.

The resulting simulation sequences, however, differ greatly from those commonly used by circuit designers during informal validation. In particular, the state of the circuit is frequently set to all X's so that any accidental sequential dependencies will be detected. During most sequences, only a small number of state or input variables are set to Boolean values, and attention is focused on the effect these values have on the output or new circuit state. Any accidental dependencies on other state or input variables will manifest themselves as X's on output or state variables that are expected to have Boolean values. Most logic simulators have not been designed for this style of simulation. Many use pessimistic methods of computing the effect of X values, causing them to produce X's where it can be shown that the true results should be Boolean values. Such a simulator provides too dull a tool for formal circuit verification, giving many false negative responses. With greater care, however, simulators can be designed to provide more accurate and efficient modeling of X's.

This methodology demonstrates several worthwhile simulation practices that could be applied even when formal verification is not sought. For example, typical simulation runs consist of many pattern sequences where the behavior of the circuit should not depend on the relative ordering of these sequences. Preceding each pattern sequence by an ERASE command would help uncover any accidental pattern sensitivities. As the memory verification example showed, a simulator can uncover more potential errors if the user can focus on small regions of the circuit at a time, setting to X those input and state variables that should not affect the behavior in this region. A common practice followed by circuit designers today is to simulate an enormous number of patterns, possibly consuming weeks of CPU time, hoping that brute force will uncover any error. By following a more disciplined methodology, shorter simulation runs could be devised that yield more reliable results.

A Design of Circuit C'

This appendix documents the design of a circuit that cannot be distinguished from a circuit that implements specification M for any simulation sequence of length less than k, as required in the proof of Theorem 1. We assume in this design that β is some sequence $[\vec{b}_1, \ldots, \vec{b}_k] \in \Phi_k$ such that $FinalOut_m(q_1, \beta) \neq FinalOut_m(q_2, \beta)$ for two states $q_1, q_2 \in Q$.

31

Let $s' = s + k$. We will define monotonic versions of the functions $next'$ and out' to be used by the simulator directly. The circuit functions are then defined by restricting the arguments to be Boolean values. Referring to Figure 5 the shift register elements in the booby trap each have two sets of inputs: a single data input t and a set of control inputs $\vec{a}$. For vector $\vec{b} \in B^p$, define the function of a shift register element $\delta_{\vec{b}} : T \times T^p \to T$ as

$$\delta_{\vec{b}}(t, \vec{a}) = \begin{cases} t, & \vec{a} = \vec{b} \\ 0, & \vec{a} \nleq \vec{b} \text{ or } t = 0 \\ X, & \vec{a} < \vec{b} \text{ and } t \neq 0 \end{cases} \tag{6}$$

That is, the input data is shifted to the output when the control inputs match those given by vector $\vec{b}$. The output is cleared to 0 when at least one control input differs from the corresponding element of $\vec{b}$ but does not equal X. To satisfy monotonicity, we adopt the convention that whenever $\vec{a} < \vec{b}$ the output equals 0 only if the data input equals 0, i.e., it does not matter whether the input is shifted or the output is cleared, and equals X otherwise.

The next state function for circuit C' is defined as

$$next'_i(\vec{z}, \vec{x}) = \begin{cases} next_i(\langle z_1, \dots z_s \rangle, \vec{x}), & i \leq s \\ \delta_{\vec{b}_1}(1, \vec{x}), & i = s + 1 \\ \delta_{\vec{b}_{i-s}}(z_{i-1}, \vec{x}) & s + 2 \leq i \leq s + k \end{cases} \tag{7}$$

The output function is defined as

$$out'_i(\vec{z}, \vec{x}) = \begin{cases} out_i(\langle z_1, \dots, z_s \rangle, \vec{x}), & i < m \\ z_{s+k} \vee out_m(\langle z_1, \dots, z_s \rangle, \vec{x}), & i = m \end{cases} \tag{8}$$

where the $\vee$ is extended monotonically as in Equation 1.

To prove that circuit C' behaves as claimed, we require the following property about the value produced by the shift register for a given input sequence.

Lemma 4 *For any $l \leq k$ and any $\alpha \in \Psi_l$,*

$$SimState'_{s+k}(\alpha) = \begin{cases} 1, & \alpha = \beta \\ 0, & \alpha \nleq \beta \\ X, & \alpha < \beta \end{cases}$$

Proof: We will show by induction on i that for any l, i such that $1 \leq l \leq i \leq k$, if we consider any sequence $\alpha \in \Psi_l$ then

$$SimState'_{s+i}(\alpha) = \begin{cases} 1, & \alpha = [\vec{b}_1, \dots \vec{b}_i] \\ 0, & \alpha \nleq [\vec{b}_1, \dots \vec{b}_i] \\ X, & \alpha < [\vec{b}_1, \dots \vec{b}_i] \end{cases} \tag{9}$$

32

The statement of the lemma then holds by letting $i = k$.

First, suppose $i = 1$, in which case either $\alpha = \epsilon$ whereby $\alpha < [\vec{b}_1]$ and $SimState'_{s-1}(\alpha) = X$, or $\alpha = [\vec{a}_1]$ for some $\vec{a}_1 \in T^p$, whereby $SimState'_{s+1}(\alpha) = \delta_{\vec{b}_1}(1, \vec{a}_1)$. Comparing equation 6 with $t = 1$ to equation 9 we see that the desired condition holds.

Now assume that equation 9 holds for some value i. Let $\alpha = \alpha' \cdot \vec{a}_{l+1}$ be a sequence where $\alpha' \in \Psi_l$ and $l \le i$. Consider the ways α can relate to the sequence $[\vec{b}_1, \ldots, \vec{b}_{i+1}]$ in Equation 9.

Equality can hold only if $l = i$ and both $\alpha' = [\vec{b}_1, \ldots, \vec{b}_i]$ and $\vec{a}_{l+1} = \vec{b}_{i+1}$. Combining equations 6, 7, and 9 for this case we get

$$SimState'_{s+i+1}(\alpha) = \delta_{\vec{b}_{i+1}}(SimState'_{s+i}(\alpha'), \vec{a}_{l+1}) = \delta_{\vec{b}_{i+1}}(1, \vec{b}_{i+1}) = 1.$$

Incomparability, i.e., $\alpha \not\le [\vec{b}_1, \ldots, \vec{b}_{i+1}]$ can hold only if either $\vec{a}_{l+1} \not\le \vec{b}_{i+1}$ or $\alpha' \not\le [\vec{b}, \ldots, \vec{b}_i]$. In the first case we have

$$SimState'_{s+i+1}(\alpha) = \delta_{\vec{b}_{i+1}}(SimState'_{s+i}(\alpha'), \vec{a}_{l+1}) = \delta_{\vec{b}_{i+1}}(t, \vec{a}_{l+1}) = 0.$$

In the second case we have

$$SimState'_{s+i+1}(\alpha) = \delta_{\vec{b}_{i+1}}(SimState'_{s+i}(\alpha'), \vec{a}_{l+1}) = \delta_{\vec{b}_{i+1}}(0, \vec{a}_{l+1}) = 0.$$

The sequences may be ordered $\alpha < [\vec{b}_1, \ldots, \vec{b}_{i+1}]$ only if either $\alpha' < [\vec{b}, \ldots, \vec{b}_i]$ and $\vec{a}_{l+1} \le \vec{b}_{i+1}$, in which case

$$SimState'_{s+i+1}(\alpha) = \delta_{\vec{b}_{i+1}}(SimState'_{s+i}(\alpha'), \vec{a}_{l+1}) = \delta_{\vec{b}_{i+1}}(X, \vec{a}_{l+1}) = X,$$

or $\alpha' = [\vec{b}, \ldots, \vec{b}_i]$ and $\vec{a}_{l+1} < \vec{b}_{i+1}$ in which case

$$SimState'_{s+i+1}(\alpha) = \delta_{\vec{b}_{i+1}}(SimState'_{s+i}(\alpha'), \vec{a}_{l+1}) = \delta_{\vec{b}_{i+1}}(1, \vec{a}_{l+1}) = X.$$

Finally, the sequences cannot be ordered $\alpha > [\vec{b}_1, \ldots, \vec{b}_{i+1}]$, because $l \le i$, and $[\vec{b}_1, \ldots, \vec{b}_{i+1}]$ is maximal.

$\square$

Lemma 5 *For any $l \le k$, any $\alpha \in \Psi_l$, and any i such that $1 \le i \le m$:*

$$SimOut'_i(\alpha) = \begin{cases} 1, & i = m \text{ and } \alpha = \beta \\ SimOut_i(\alpha), & \text{else} \end{cases}$$

Proof: For all cases except where $i = m$ and $\alpha < \beta$, this result follows from the definition of *out'* and from Lemma 4. For $\alpha < \beta$, Lemma 4 tells us that the shift register output

will equal X. However, given that circuit C implements specification $\mathcal{M}$, Lemma 3 shows that both $SimOut_m(\alpha) \leq FinalOut_m(q_1, \beta)$ and $SimOut_m(\alpha) \leq FinalOut_m(q_2, \beta)$. Since these two values are unequal, we must have

$$SimOut_m(\alpha) = SimOut'_m(\alpha) = X.$$

$\square$

References

[1] Barrow, H. G. Proving the correctness of digital hardware designs. *VLSI Design V*, 7 (July 1984), 64–77.

[2] Barrow, H. G. VERIFY: a program for proving correctness of digital hardware designs. *Artificial Intelligence 24* (1984), 437–491.

[3] Browne, M. C., Clarke, E. M., Dill, D. L., and Mishra, B. Automatic verification of sequential circuits using temporal logic. *IEEE Transactions on Computers C-35*, 12 (Dec. 1986), 1035–1044.

[4] Bryant, R. E. A switch-level model and simulator for MOS digital systems. *IEEE Transactions on Computers C-33*, 2 (Feb. 1984), 160–177.

[5] Bryant, R. E. Symbolic verification of MOS circuits. *1985 Chapel Hill Conference on VLSI*, Fuchs, H., Ed. Computer Science Press, Rockville, MD, 1985, 419–438.

[6] Brzozowski, J. A., and Yoeli, M. On a ternary model of gate networks. *IEEE Transactions on Computers C-28*, 3 (March 1979), 178–183.

[7] Dill, D. L., and Clarke, E. M. Automatic verification of asynchronous circuits using temporal logic. *IEE Proceedings 133*, Pt. E, 5 (Sept. 1986), 276–282.

[8] Floyd, R. W. "Assigning meanings to programs," *Proc. Symp. in Applied Mathematics, 19—Mathematical Aspects of Computer Science*, Schwartz, J. T., Ed. AMS, 1967, 19–32.

[9] German, S. M., and Wang, Y. Formal verification of parameterized hardware designs. *Int. Conf. on Computer Design*, IEEE, 1985, 549–552.

[10] Glasser, L. A., and Dobberpuhl, D. W. *The Design and Analysis of VLSI Circuits*, Addison-Wesley, Reading, MA, 1985.

[11] Hoare, C. A. R. An axiomatic basis for computer programming. *Comm. ACM 12* (1969), 576–580.

[12] Jephson, J. S., McQuarrie, R. P., and Vogelsberg, R. E. A three-level design verification system. *IBM Systems Journal 8*, 3 (1969), 178–188.

[13] Kleene, S. C. Representation of events in nerve nets and finite automata. *Automata Studies*, Shannon, C. E., and McCarthy, J., Ed. Princeton University Press, Princeton, NJ, 1956, 3–41.

[14] Kohavi, Z. *Switching and Finite Automata Theory*, McGraw-Hill, New York, 1970.

[15] Lengauer, T., and Näher, S. An analysis of ternary simulation as a tool for race detection. *Integration, the VLSI Journal 4*, 4 (Dec. 1986), 309–330.

[16] Milne, G. J. CIRCAL: a calculus for circuit description. *Integration 1*, 2&3 (Oct. 1983) 121–160.

[17] Milne, G. J. A model for hardware description and verification. *21st Design Automation Conference*, ACM and IEEE, 1984.

[18] Moore, E. F. Gedanken-experiments on sequential machines. *Automata Studies*, Shannon C. E., and McCarthy, J., Ed. Princeton University Press, Princeton, NJ, 1956, 129–153.

[19] Perles M., Rabin M. O., and Shamir E. The theory of definite automata. *IEEE Transactions on Electronic Computers EC-12*, 6 (June, 1963), 233–243.

[20] Roth, J. P. *Computer Logic, Testing, and Verification*, Computer Science Press, Rockville, MD, 1980.

[21] Shostak, R. E. Verification of VLSI designs. *Proceedings of the Third Caltech Conference on VLSI*, Bryant, R., Ed. Computer Science Press, Rockville, MD, 1983, 185–206.

[22] Wagner, T. J. *Hardware Verification*. Ph.D. Thesis, Dept. Comp. Sci., Stanford Univ., 1977.

[23] Weise, D. *Automatic Formal Verification of Synchronous MOS VLSI Designs*. Ph.D. Thesis, Dept. Elec. Eng. and Comp. Sci., Massachusetts Inst. of Tech., 1986.

[24] Weise, D. Verifying MOS circuits, *24th Design Automation Conference*, ACM and IEEE, 1987.

[25] Yoeli, M., and Rinon, S. Application of ternary algebra to the study of static hazards, *J. ACM 11*, 1 (Jan. 1964), 84–97.

END

DATE
FILMD
3 — 88
DTIC